

WHAT IS CLAIMED IS:

1 1. A method for processing data in an application data store comprising:
2 producing at least a first snapshot of an application data store, the application
3 data store configured to receive data by way of write operations issued from a host device;
4 producing a journal entry for each write operation issued from the host device;
5 storing each journal entry in a journal data store, thereby accumulating a list of
6 journal entries;
7 monitoring an amount of free space on the journal data store; and
8 when the free space falls below a first threshold value, then removing one or
9 more journal entries from the journal data store, thereby increasing the free space, wherein
10 enough of the journal entries are removed so that the free space rises above a second
11 threshold value.

1 2. The method of claim 1 wherein removing one or more journal entries
2 includes updating the first snapshot by applying one or more journal entries to the first
3 snapshot, beginning with an oldest journal entry, wherein journal entries applied to the first
4 snapshot are removed from the list of journal entries thereby increasing the free space of the
5 journal data store.

1 3. The method of claim 1 wherein removing one or more journal entries
2 includes identifying a selected snapshot that is earlier in time than an oldest journal entry and
3 is closest in time to the oldest journal entry than other snapshots and updating the selected
4 snapshot with one or more journal entries beginning with the oldest journal entry.

1 4. The method of claim 1 wherein removing one or more journal entries
2 includes looking for a most recent snapshot, removing journal entries in the journal data store
3 that are earlier in time than the most recent snapshot.

1 5. The method of claim 1 wherein the first threshold value and the second
2 threshold value are different.

1 6. The method of claim 1 wherein each journal entry comprises a fixed-
2 size header portion and a variable-size data portion, wherein the journal data store comprises
3 a first storage area within which a plurality of header portions are defined and a second
4 storage area for storing a plurality of data portions, wherein producing a journal entry

5 includes allocating one of the journal headers and allocating a sufficient amount of space in
6 the data portion to contain data associated with the write operation.

1 7. The method of claim 6 wherein storing each journal entry includes
2 allocating one of the journal headers from the first portion such that the journal headers are
3 always allocated in sequential manner and when the last journal header is allocated then
4 allocating journal headers from the beginning of the first portion in sequential manner.

1 8. A data processing method comprising:
2 producing at least a first snapshot of at least a portion of an application data
3 store, the application data store being configured to receive write operations issued from a
4 host device;
5 recording a plurality of journal entries, each journal entry being recorded for a
6 write operation issued from the host device, each journal entry being stored in a journal data
7 store thereby consuming an amount of free space of the journal data store; and
8 updating the at least first snapshot with at least one journal entry, including
9 associating space consumed by the at least one journal entry with the free space thereby
10 increasing the free space of the journal data store.

1 9. The method of claim 8 wherein the step of updating includes
2 periodically monitoring the free space of the journal data store and if the free space falls
3 below a first threshold then updating the at least first snapshot.

1 10. The method of claim 9 wherein the step of updating the at least first
2 snapshot is repeated for a number of journal entries so that the free space rises above a
3 second threshold value.

1 11. The method of claim 8 further comprising:
2 associating sequence numbers to the first snapshot and to the journal entries;
3 receiving a target time;
4 determining a beginning journal entry based on the sequence numbers
5 associated with the journal entries and the sequence number associated with the first
6 snapshot;
7 updating the first snapshot by applying the beginning journal entry to it; and
8 performing additional updates to the first snapshot by applying journal entries
9 that are subsequent in time to the beginning journal entry and prior in time to the target time.

1 12. The method of claim 8 further comprising:
2 producing additional snapshots thereby accumulating a plurality of snapshots;
3 associating sequence numbers to the snapshots and to the journal entries;
4 receiving a target time;
5 determining a selected snapshot based on the target time;
6 determining a beginning journal entry based on the sequence numbers
7 associated with the journal entries and the sequence number associated with the selected
8 snapshot;
9 updating the selected snapshot by applying the beginning journal entry to it;
10 and
11 performing additional updates to the selected snapshot by applying journal
12 entries that are subsequent in time to the beginning journal entry and prior in time to the
13 target time.

1 13. The method of claim 12 wherein the selected snapshot is a first
2 snapshot.

1 14. The method of claim 12 wherein the selected snapshot is closest in
2 time to the target time and prior in time to the target time.

1 15. A method for processing data comprising:
2 producing a plurality of snapshots of one or more portions of an application
3 data store, the data volume configured to receive data by way of write operations issued from
4 a host device;
5 recording a journal entry for each write operation issued from the host device,
6 each journal entry being stored in a journal data store thereby accumulating a list of journal
7 entries and consuming free space on the journal data store;
8 monitoring a remaining free space on the journal data store; and
9 when the remaining free space falls below a first threshold value, then
10 removing one or more journal entries by identifying a selected snapshot and removing journal
11 entries earlier in time than the selected snapshot.

1 16. The method of claim 15 wherein the selected snapshot is at least a
2 predetermined length of time later than an oldest journal entry,

1 17. The method of claim 15 further comprising associating sequence
2 numbers to the snapshots and to the journal entries, wherein the selected snapshot is
3 associated with a sequence number that is greater than a sequence number associated with an
4 oldest journal entry by a predetermined amount.

1 18. A method for processing data comprising:
2 producing a plurality of snapshots of an application data store, the data volume
3 configured to receive data by way of write operations issued from a host device;
4 producing a journal entry for each write operation issued from the host device;
5 storing each journal entry in a journal data store, thereby accumulating a list of
6 journal entries;
7 monitoring an amount of free space on the journal data store; and
8 when the free space falls below a first threshold value, then:
9 identifying a selected snapshot that is earlier in time than an oldest
10 journal entry and is closest in time to the oldest journal entry than other snapshots;
11 updating the selected snapshot with one or more journal entries
12 beginning with the oldest journal entry, thereby reducing the list of journal entries;
13 and
14 repeating the updating step for additional journal entries until the free
15 space rises above a second threshold value.

1 19. The method of claim 18 further comprising associating sequence
2 numbers to the snapshots and to the journal entries, wherein the selected snapshot is
3 determined based on the sequence numbers of the snapshots and the sequence number of the
4 oldest journal entry.

1 20. The method of claim 18 wherein journal entries are sequentially
2 allocated and contiguously arranged in the journal data store in a manner representative of a
3 sequential list of journal entries, wherein the step of storing includes wrapping to a beginning
4 of the list of journal entries to reuse earlier journal entries when an end of the list of journal
5 entries is reached.

1 21. A method for processing data comprising:
2 producing a plurality of snapshots of an application data store, the application
3 data store configured to receive data by way of write operations issued from a host device;

4 recording a journal entry for each write operation issued from the host device,
5 the journal entry being stored in a journal data store, thereby accumulating a list of journal
6 entries in the journal data store;
7 producing a plurality of sequential numbers;
8 for each snapshot and for each journal entry, associating one of the numbers
9 therewith;
10 monitoring an amount of free space on the journal data store; and
11 when the free space falls below a first threshold value, then:
12 identifying a selected snapshot by comparing sequence numbers
13 associated with the snapshots against a sequence number associated with an oldest
14 journal entry;
15 updating the selected snapshot with one or more journal entries
16 beginning with the oldest journal entry; and
17 repeating the updating step with additional journal entries until the free
18 space rises above a second threshold value.

1 22. The method of claim 21 wherein the step of recording is initiated prior
2 to producing a snapshot, so each write operation that issues during a period of time that a first
3 snapshot is being produced will be recorded in a journal entry.

1 23. The method of claim 21 wherein the selected snapshot is earlier in time
2 than the oldest journal entry.

1 24. The method of claim 21 wherein the selected snapshot is earlier in time
2 than the oldest journal entry and closest in time to the oldest journal entry than other
3 snapshots.

1 25. The method of claim 21 wherein the first threshold is different from
2 the second threshold.

1 26. A method for processing data in a data storage system comprising:
2 initiating journal processing wherein a journal entry is created for each write
3 operation that is sent to the data storage system, thus producing a plurality of journal entries,
4 the journal entries being stored in a backup storage system;
5 subsequent to the step of initiating journal processing, taking one or more
6 snapshots of a portion of the data contained in the data storage system;
7 associating a sequence number to each journal entry when it is created and to
8 each snapshot when it is taken, wherein sequence numbers associated with the journal entries
9 and with the snapshots are part of the same sequence; and
10 determining whether to apply some of the journal entries to one of the
11 snapshots in order to reduce the amount of space in the backup storage system that is
12 consumed by the journal entries, based on the amount of available storage capacity in the
13 backup storage system.

1 27. The method of claim 26 wherein if a determination is made to apply
2 some of the journal entries to one of the snapshots, then:
3 identifying an oldest journal entry;
4 identifying a selected snapshot based on a sequence number associated with
5 the oldest journal entry and on sequence numbers associated with the snapshots; and
6 updating the selected snapshot with one or more journal entries, beginning
7 with the oldest journal entry.

1 28. A storage system for processing data comprising:
2 a production data store configured to receive write operations from a host
3 device;
4 a snapshot data store configured to store one or more snapshots of at least a
5 portion of the production data store;
6 a journal data store configured to store one or more journal entries; and
7 a controller configured to:
8 access the production data store and the snapshot data store to store
9 one or more snapshots of at least a portion of the production data store on the
10 snapshot data store;
11 access the journal data store to record a journal entry for each write
12 operation from the host device; and

13 update one of the snapshots with some, but not all, of the journal
14 entries thereby leaving some journal entries for a data recovery operation.

1 29. The storage system of claim 28 wherein the controller is further
2 configured to:
3 associate sequence numbers to the snapshots and to the journal entries;
4 receive a target time;
5 select a snapshot based on the target time;
6 determine a beginning journal entry based on the sequence numbers associated
7 with the journal entries and the sequence number of the selected snapshot;
8 update the selected snapshot by applying the beginning journal entry to it; and
9 perform additional updates to the selected snapshot by applying journal entries
10 that are later in time than the beginning journal entry and earlier in time than the target time.

1 30. A storage system for processing data comprising:
2 means for receiving communications from a host device, the communications
3 including one or more write operations;
4 first storage means for storing production data;
5 second storage means for storing snapshots;
6 third storage means for storing journal entries;
7 control means for:
8 storing onto the second storage means one or more snapshots of at least
9 a portion of the production data;
10 recording onto the third storage means a journal entry for each write
11 operation thereby producing a plurality of journal entries and reducing an amount of
12 free space on the third storage means; and
13 updating one of the snapshots according to one or more journal entries
14 thereby increasing the amount of free space on the third storage means.

1 31. The storage system of claim 30 wherein the updating can be performed
2 periodically or at intervals based on the amount of free space.